



Ivan Pepelnjak,
CCIE™ #1354

Jim Guichard,
CCIE #2069

MPLS AND VPN ARCHITECTURES

**A practical guide to understanding, designing and
deploying MPLS and MPLS-enabled VPNs**

CISCO SYSTEMS



CISCO PRESS

ciscopress.com

MPLS and VPN Architectures

Jim Guichard, CCIE #2069

Ivan Pepelnjak, CCIE #1354



Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

MPLS and VPN Architectures

Jim Guichard, CCIE #2069

Ivan Pepelnjak, CCIE #1354

Copyright© 2001 Cisco Press

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:

Cisco Press

201 West 103rd Street

Indianapolis, IN 46290 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 3 4 5 6 7 8 9 0 03 02 01

3rd Printing March 2001

Library of Congress Cataloging-in-Publication Number: 00-105168

ISBN: 1-58705-002-1

Warning and Disclaimer

This book is designed to provide information about Multiprotocol Label Switching (MPLS) and Virtual Private Networks (VPN). Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The author, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors and are not necessarily those of Cisco Systems, Inc.

Feedback Information


At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through e-mail at ciscopress@mcp.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.



Contents at a Glance

Part I	MPLS Technology and Configuration 3
Chapter 1	Multiprotocol Label Switching (MPLS) Architecture Overview 5
Chapter 2	Frame-mode MPLS Operation 23
Chapter 3	Cell-mode MPLS Operation 49
Chapter 4	Running Frame-mode MPLS Across Switched WAN Media 65
Chapter 5	Advanced MPLS Topics 73
Chapter 6	MPLS Migration and Configuration Case Study 97
Part II	MPLS-based Virtual Private Networks 113
Chapter 7	Virtual Private Network (VPN) Implementation Options 115
Chapter 8	MPLS/VPN Architecture Overview 145
Chapter 9	MPLS/VPN Architecture Operation 167
Chapter 10	Provider Edge (PE) to Customer Edge (CE) Connectivity Options 207
Chapter 11	Advanced MPLS/VPN Topologies 237
Chapter 12	Advanced MPLS/VPN Topics 249
Chapter 13	Guidelines for the Deployment of MPLS/VPN 319
Chapter 14	Carrier's Carrier and Inter-provider VPN Solutions 357
Chapter 15	IP Tunneling to MPLS/VPN Migration Case Study 387
Appendix A	Tag-Switching and MPLS Command Reference 405
Index	408

Contents

Part I	MPLS Technology and Configuration 3
Chapter 1	Multiprotocol Label Switching (MPLS) Architecture Overview 5
	Scalability and Flexibility of IP-based Forwarding 5
	Network Layer Routing Paradigm 5
	Differentiated Packet Servicing 9
	Independent Forwarding and Control 10
	External Routing Information Propagation 10
	Multiprotocol Label Switching (MPLS) Introduction 11
	MPLS Architecture—The Building Blocks 13
	Label Imposition at the Network Edge 15
	MPLS Packet Forwarding and Label Switched Paths 17
	Other MPLS Applications 19
	Summary 20
Chapter 2	Frame-mode MPLS Operation 23
	Frame-mode MPLS Data Plane Operation 25
	MPLS Label Stack Header 26
	Label Switching in Frame-mode MPLS 28
	Label Bindings and Propagation in Frame-mode MPLS 31
	LDP/TDP Session Establishment 31
	Label Binding and Distribution 34
	Convergence in a Frame-mode MPLS Network 37
	Penultimate Hop Popping 40
	MPLS Interaction with the Border Gateway Protocol 42
	Summary 45
Chapter 3	Cell-mode MPLS Operation 49
	Control-plane Connectivity Across an LC-ATM Interface 51
	MPLS Control-plane Connectivity in Cisco IOS Software 52
	Control-plane Implementation in an ATM Switch 54
	Labeled Packet Forwarding Across an ATM-LSR Domain 55

	Label Allocation and Distribution Across an ATM-LSR Domain	56
	VC Merge	58
	Convergence Across an ATM-LSR Domain	61
	Summary	62
Chapter 4	Running Frame-mode MPLS Across Switched WAN Media	65
	Frame-mode MPLS Operation Across Frame Relay	65
	Frame-mode MPLS Operation Across ATM PVCs	67
	Frame-mode and Cell-mode MPLS Across the Same ATM Interface	69
	Summary	70
Chapter 5	Advanced MPLS Topics	73
	Controlling the Distribution of Label Mappings	73
	MPLS Encapsulation Across Ethernet Links	77
	IP MTU Path Discovery	78
	Ethernet Switches and MPLS MTU	80
	MPLS Loop Detection and Prevention	81
	Loop Detection and Prevention in Frame-mode MPLS	81
	Loop Detection and Prevention in Cell-mode MPLS	82
	Traceroute Across an MPLS-enabled Network	88
	Route Summarization Within an MPLS-enabled Network	92
	Summary	93
Chapter 6	MPLS Migration and Configuration Case Study	97
	Migration of the Backbone to a Frame-mode MPLS Solution	97
	Pre-migration Infrastructure Checks	99
	Cisco Express Forwarding (CEF) Requirements	100
	Addressing the Internal BGP Structure	101
	Migration of Internal Links to MPLS	103
	Removal of Unnecessary BGP Peering Sessions	105
	Migration of an ATM-based Backbone to Frame-mode MPLS	106
	Cell-mode MPLS Migration	108
	Summary	110

Part II MPLS-based Virtual Private Networks 113**Chapter 7 Virtual Private Network (VPN) Implementation Options 115**

- Virtual Private Network Evolution 115
 - Modern Virtual Private Networks 118
- Business Problem-based VPN Classification 118
- Overlay and Peer-to-peer VPN Model 121
 - Overlay VPN Model 121
 - Peer-to-peer VPN Model 123
- Typical VPN Network Topologies 129
 - Hub-and-spoke Topology 129
 - Partial- or Full-mesh Topology 133
 - Hybrid Topology 134
 - Simple Extranet Topology 135
 - Central-services Extranet 137
 - VPDN Topology 139
 - Managed Network VPN Topology 141
- Summary 142

Chapter 8 MPLS/VPN Architecture Overview 145

- Case Study: Virtual Private Networks in SuperCom Service Provider Network 146
- VPN Routing and Forwarding Tables 149
- Overlapping Virtual Private Networks 151
- Route Targets 155
- Propagation of VPN Routing Information in the Provider Network 156
 - Multiprotocol BGP in the SuperCom Network 158
- VPN Packet Forwarding 161
- Summary 163

Chapter 9 MPLS/VPN Architecture Operation 167

- Case Study: Basic MPLS/VPN Intranet Service 168
- Configuration of VRFs 170
- Route Distinguishers and VPN-IPv4 Address Prefixes 171
 - Configuration of the Route Distinguisher 175

- BGP Extended Community Attribute 177
 - Route Target BGP Extended Community 178
 - Site of Origin BGP Extended Community 180
 - BGP Extended Community Attribute Format 182
- Basic PE to CE Link Configuration 184
 - PE to CE Link Configuration—Static Routing 184
 - PE to CE Link Configuration—RIP Version 2 185
- Association of Interfaces to VRFs 187
- Multiprotocol BGP Usage and Deployment 188
 - Configuration of Multiprotocol BGP 191
 - Enhanced BGP Decision Process for VPN-IPv4 Prefixes 194
- Outbound Route Filtering (ORF) and Route Refresh Features 196
 - Automatic Route Filtering on PE-routers 197
 - Refreshing Routing Information Between PE-routers 199
 - ORF for PE-routers 201
- MPLS/VPN Data Plane—Packet Forwarding 203
- Summary 205

Chapter 10 Provider Edge (PE) to Customer Edge (CE) Connectivity Options 207

- VPN Customer Access into the MPLS/VPN Backbone 207
- BGP-4 Between Service Provider and Customer Networks 209
- Open Shortest Path First (OSPF) Between PE- and CE-routers 212
- Separation of VPN Customer Routing Information 214
- Propagation of OSPF Routes Across the MPLS/VPN Backbone 217
 - BGP Extended Community Attribute for OSPF Routes 219
- PE-to-CE Connectivity—OSPF with Site Area 0 Support 220
- PE-to-CE Connectivity—OSPF Without Site Area 0 Support 224
- VPN Customer Connectivity—MPLS/VPN Design Choices 226
 - Migrating Customers Using iBGP in Their Network to MPLS/VPN Service 230
 - Autonomous System Number Override 232
- Summary 234

Chapter 11 Advanced MPLS/VPN Topologies 237

- Intranet and Extranet Integration 237
- Central Services Topology 240

MPLS/VPN Hub-and-spoke Topology 242
Deployment of the AllowAS-in Feature 245

Summary 247

Chapter 12 Advanced MPLS/VPN Topics 249

MPLS/VPN: Scaling the Solution 251

Routing Convergence Within an MPLS-enabled VPN Network 251
Convergence Within the Service Provider Backbone 253
Convergence Between VPN Sites 255

Advertisement of Routes Across the Backbone 260
BGP for VPN-IPv4 and IPv4 Routing Information 261
Full Mesh of MP-iBGP Between PE-routers 265
Separation of MP-iBGP Sessions Between PE-routers 266

Introduction of Route Reflector Hierarchy 268
Route Reflection of PE Routes to Aid Scaling 270
Route Reflector Partitioning 272
Standard Community Filtering on PE-routers 274
Route Target Attribute-based Filtering on Route Reflectors 277
Route Reflection and ORF Capability 279

BGP Confederations Deployment 280
BGP Confederations—Single IGP Environment 285
BGP Confederations—Multiple IGP Environment 286

PE-router Provisioning and Scaling 291

Additional Connectivity Requirements—Internet Access 292

Internet Connectivity Through Firewalls 293

Internet Access—Static Default Routing 295

Separate BGP Session Between PE- and CE-routers 300

Internet Connectivity Through Dynamic Default Routing 308
Dynamic Default Routing—Route Target Assignment 308
Association of the Global Routing Table with a VRF 310

Additional Lookup in the Global Routing Table 314

Internet Connectivity Through a Different Service Provider 315

Summary 316

Chapter 13	Guidelines for the Deployment of MPLS/VPN 319
	Introduction to MPLS/VPN Deployment 319
	IGP to BGP Migration of Customer Routes 319
	Multiprotocol BGP Deployment in an MPLS/VPN Backbone 324
	VPN Routes and Next-hop Forwarding 325
	PE-router Loopback Address Configuration 328
	MPLS/VPN Deployment on LAN Interfaces 335
	Network Management of Customer Links 338
	Advertisement of Routes with Different Extended Communities 339
	Use of Standard BGP Communities for Route Filtering 343
	Advertisement of Routes with Different Route Targets Using Export Maps 348
	Use of Traceroute Across an MPLS/VPN Backbone 351
	Summary 354
Chapter 14	Carrier's Carrier and Inter-provider VPN Solutions 357
	Carrier's Carrier Solution Overview 358
	Carrier's Carrier Architecture—Topologies 361
	ISP with No MPLS Deployment Within POP Sites 362
	ISP with MPLS Deployed Within POP Sites 368
	Hierarchical Virtual Private Networks 372
	Inter-provider VPN Solutions 374
	Inter-provider VPN—Exchange of VPN-IPv4 Across Boundaries 376
	Inter-provider VPN—Multi-hop eBGP Between Customer Sites 382
	Summary 384
Chapter 15	IP Tunneling to MPLS/VPN Migration Case Study 387
	Existing VPN Solution Deployment—IP Tunneling 388
	Definition of VPNs and Routing Policies for PE-routers 390
	Definition of VRFs Within the Backbone Network 391
	VRF and Routing Policies for SampleNet VPN Sites 392
	VRF and Routing Policies for SampleNet Internet Access 393
	VRF and Routing Policies for Internet Access Customers 394
	MPLS/VPN Migration—Staging and Execution 394
	Migration of the SampleNet Central Site 395

Configuration of MP-iBGP on BGP Route Reflectors 398

Configuration of MP-iBGP on TransitNet PE-routers 400

Migration of VPN Sites onto the MPLS/VPN Solution 401

Summary 401

Appendix A Tag-Switching and MPLS Command Reference 405

Index 408

Virtual Private Network (VPN) Implementation Options

A Virtual Private Network (VPN) is defined loosely as a network in which customer connectivity amongst multiple sites is deployed on a shared infrastructure with the same access or security policies as a private network. With the recent advent of marketing activities surrounding the term VPNs, from new technologies supporting VPNs to a flurry of VPN-enabled products and services, you might think that the VPN concept is a major technology throughput. However, as is often the case, VPN is a concept that is more than 10-years old and is well known in the service provider market space.


The new technologies and products merely enable more reliable, scalable, and more cost-effective implementation of the same product. With the cost reduction and enhanced scalability associated with new VPN technologies, it's not surprising that VPN services are among the major drivers for Multiprotocol Label Switching (MPLS) deployment in service provider and enterprise networks.

Before discussing a technology (VPN services based on MPLS) designed to solve a problem (cost-effective VPN implementation), it's always advantageous to focus on the problem first, which is what we do in this chapter.

This chapter gives you an overview of VPN services, common VPN terminology, and detailed classification of various VPN usages and topologies that are encountered most often. This chapter also provides an overview of technologies that were used traditionally to implement Virtual Private Networks either on individual service provider backbones or over the public Internet.

Virtual Private Network Evolution

Initial computer networks were implemented with two major technologies: *leased lines* for permanent connectivity and *dial-up lines* for occasional connectivity requirements. Figure 7-1 shows a typical network from those days.

Figure 7-1 Typical Computer Network from 15 Years Ago


The initial computer network implementation provided the customers with good security (capturing data off leased lines requires dedicated equipment and physical access to the wires), but did not provide cost-effective implementation due to two reasons:

- The typical traffic profile between any two sites in a network varies based on the time of day, the day of the month, and even the season (for example, traffic at retail stores increases around Christmas season).
- The end-users always request fast responses, resulting in a high bandwidth requirement between sites, but the dedicated bandwidth available on the leased lines is used only part of the time (when the users are active).

These two reasons prompted the data communication industry and service providers to develop and implement a number of statistical multiplexing schemas that provided the customers with a service that was almost an equivalent to leased lines. This service was cheaper, however, due to the statistical benefits the service provider could achieve from a large customer base. The first *virtual* private networks were based on such technologies as X.25 and Frame Relay, and, later, SMDS and ATM. Figure 7-2 shows a typical VPN built with these technologies (for example, Frame Relay).

As you can see in Figure 7-2, the overall VPN solution has a number of components:

- The service provider is the organization that owns the infrastructure (the equipment and the transmission media) that provides emulated leased lines to its customers. The service provider in this scenario offers a customer a *Virtual Private Network Service*.

Figure 7-2 Typical Frame Relay Network

- The customer connects to the service provider network through a *Customer Premises Equipment (CPE)* device. The CPE is usually a Packet Assembly and Disassembly (PAD) device that provides plain terminal connectivity, a bridge, or a router. The CPE device is also sometimes called a *Customer Edge (CE)* device.
- The CPE device is connected through transmission media (usually a leased line, but could also be a dial-up connection) to the service provider equipment, which could be an X.25, Frame Relay, or ATM switch, or even an IP router. The edge service provider device is sometimes called the *Provider Edge (PE)* device.
- The service provider usually has additional equipment in the core of the service provider network (also called the *P-network*). These devices are called *P-devices* (for example, P-switches or P-routers).
- A contiguous part of the customer network is called a *site*. A site can connect to the P-network through one or several transmission lines, using one or several CPE and PE devices, based on the redundancy requirements.
- The emulated leased line provided to the customer by the service provider in the overlay VPN model (see the section, "Overlay and Peer-to-peer VPN Model," later in this chapter for more details) frequently is called a *Virtual Circuit (VC)*. The VC can be either constantly available (*Permanent Virtual Circuit [PVC]*) or established on demand (*Switched Virtual Circuit [SVC]*). Some technologies used special terms for VCs, for example Data Link Connection Identifier (DLCI) in Frame Relay.
- The service provider can charge either a flat rate for the VPN service, which normally depends on the bandwidth available to the customer, or a usage-based rate, which can depend on the volume of data exchanged or the duration of data exchange.

Modern Virtual Private Networks

With the introduction of new technologies in the service provider networks and new customer requirements, the VPN concept became more and more complex. Vendors introduced different and often conflicting terms, which further increased the complexity. The modern VPN services thus can span a variety of technologies and topologies. The only way to cope with this diversity is to introduce VPN classification, which you can do using four criteria:

- The business problem a VPN is trying to solve. The major classes of business problems are intracompany communication (lately, also called *intranet*), inter-company communication (also called *extranet*), and access for mobile users (also called *Virtual Private Dialup Network*).
- The OSI layer at which the service provider exchanges the topology information with the customer. Major categories here are the *overlay model*, where the service provider provides the customer with only a set of point-to-point (or multipoint) links between the customer sites, and the *peer model*, where the service provider and the customer exchange Layer 3 routing information.
- The Layer 2 or Layer 3 technology used to implement the VPN service within the service provider network, which can be X.25, Frame Relay, SMDS, ATM, or IP.
- The topology of the network, which can range from simple hub-and-spoke topology to fully meshed networks and multilevel hierarchical topologies in larger networks.

Business Problem-based VPN Classification

The three business problems a typical organization is trying to solve with a Virtual Private Network are:

- Intra-organizational communication (*intranet*).
- Communication with other organizations (*extranet*).
- Access of mobile users, home workers, remote office, and so on, through inexpensive dial-up media (*Virtual Private Dial-up Network*)


The three types of VPN solutions usually span most of the topologies and technologies offered by VPN service providers, but differ greatly in the level of security required in their implementation.

Intra-organizational communications usually are not protected well by the end hosts or the firewalls. The VPN service used to implement intra-organizational communication therefore must offer high levels of isolation and security. Intra-organizational communications also require guaranteed quality of service for mission-critical processes.


These are the two major reasons why we don't see many organizations using Internet, which cannot offer end-to-end quality of service, isolation, or security, as the infrastructure for their intra-organizational communications. Intranet VPNs were thus usually implemented with traditional technologies like X.25, Frame Relay, or ATM.

Inter-organizational communications frequently take place between central sites of the organizations—usually using dedicated security devices, such as firewalls or encryption gear similar to the setup demonstrated in Figure 7-3. These communications also might have less stringent quality of service requirements. This set of requirements makes the Internet more and more suitable for inter-organizational communications; therefore, it's no surprise that more and more business-to-business traffic takes place over the Internet.

Figure 7-3 *Typical Extranet Setup*



Remote user access into a corporate network, typically from changing or unknown locations, is always riddled with security issues, which have to be resolved on an end-to-end basis using such technologies as encryption or one-time passwords. Thus, the security requirements for VPDN services were never as high as the requirements for Intranet communications. It's no surprise that most of the VPDN services today are implemented on top of Internet Protocol (IP), either over the Internet or using the private backbone of a service provider, as illustrated in Figure 7-4. The protocols used to implement VPDN service over IP include Layer 2 Forwarding (L2F) or Layer 2 Transport Protocol (L2TP).

Figure 7-4 Service Provider Offering Separate VPDN Backbone

The VPDN technology uses a number of special terms that are unique to the VPDN world:

- **Network Access Server (NAS)**—The Remote Access Server (RAS) managed by the service provider that accepts the customer call, performs the initial authentication, and forwards the call (via L2F or L2TP) to the customer's gateway.
- **Home Gateway**—A customer-managed router that accepts the call forwarded by the NAS, performs additional authentication and authorization, and terminates the PPP session from the dial-up user. The PPP session parameters (including network addresses, such as an IP address) are negotiated between the dial-up user and the home gateway; NAS only forwards frames of Point-to-Point Protocol (PPP) between the two.

NOTE

The details of VPDN, L2F, and L2TP are beyond the scope of this book. Please refer to *RFC 2341 Cisco Layer Two Forwarding (Protocol) "L2F"* and *RFC 2661 Layer Two Tunneling Protocol "L2TP"* for additional information on these topics.

Overlay and Peer-to-peer VPN Model

Two VPN implementation models have gained widespread use:

- The overlay model, where the service provider provides emulated leased lines to the customer.
- The peer-to-peer model, where the service provider and the customer exchange Layer 3 routing information and the provider relays the data between the customer sites on the optimum path between the sites and without the customer's involvement.

NOTE


One might argue that the case where the customer and the provider use the same Layer 2 technology (for example, Frame Relay or ATM switches) also constitutes a peer-to-peer model, but because we focus on Layer 3 VPN services here, we will not consider this scenario. Similarly, a humorous person might call a leased line service a Layer 1 peer-to-peer model.

Overlay VPN Model

The overlay VPN model is the easiest to understand because it provides very clear separation between the customer's and the service provider's responsibilities:


- The service provider provides the customer with a set of emulated leased lines. These leased lines are called VCs, which can be either constantly available (PVCs) or established on demand (SVCs). Figure 7-5 shows the topology of a sample overlay VPN and the VCs used in it.

Figure 7-5 Sample Overlay VPN Network



- The customer establishes router-to-router communication between the Customer Premises Equipment (CPE) devices over the VCs provisioned by the service provider. The routing protocol data is always exchanged between the customer devices, and the service provider has no knowledge of the internal structure of the customer network. Figure 7-6 shows the routing topology of the VPN network in Figure 7-5.

Figure 7-6 *Routing in Sample Overlay VPN Network*



The QoS guarantees in the overlay VPN model usually are expressed in terms of bandwidth guaranteed on a certain VC (Committed Information Rate or CIR) and maximum bandwidth available on a certain VC (Peak Information Rate or PIR). The committed bandwidth guarantee usually is provided through the statistical nature of the Layer 2 service but depends on the overbooking strategy of the service provider. This means that the committed rate is not actually guaranteed although the provider can provision a Minimum Information Rate (MIR) that effectively is nailed up across the Layer 2 infrastructure.

NOTE

The committed bandwidth guarantee is also only a guarantee of the bandwidth between two points in the customer network. Without a full traffic matrix for all traffic classes, it's hard for the customer to engineer guarantees in most overlay networks. It's also hard to provide multiple classes of service because the service provider cannot differentiate the traffic in the middle of the network. Working around this by creating multiple connections (for example, Frame Relay PVCs) between the customer sites only increases the overall cost of the network.

Overlay VPN networks can be implemented with a number of switched WAN Layer 2 technologies, including X.25, Frame Relay, ATM, or SMDS. In the last years, overlay VPN networks also have been implemented with IP-over-IP tunneling, both in private IP backbones and over the public Internet. The two most commonly used IP-over-IP tunneling methods are Generic Route Encapsulation (GRE) tunneling and IP Security (IPSec) encryption.

NOTE

This book does not discuss the various Layer 2 and Layer 3 overlay VPN technologies in detail because they are covered well in other Cisco Press publications and are beyond the scope of this book. For more information on Layer 2 WAN technologies, please refer to *Internetworking Technologies Handbook*, Second Edition, from Cisco Press (ISBN 1-57870-102-3). For a description of IP-over-IP tunneling and IPSec encryption, please see *RFC 1702 – Generic Routing Encapsulation over IPv4 networks*, *RFC 2401 – Security Architecture for the Internet Protocol*, and *Enhanced IP Services for Cisco Networks* from Cisco Press (ISBN 1-57870-106-6).

Although it's relatively easy to understand and implement, the overlay VPN model nevertheless has a number of drawbacks:

- It's well suited to non-redundant configurations with a few central sites and many remote sites, but becomes exceedingly hard to manage in a more meshed configuration (see also the section, "Typical VPN Network Topologies," later in this chapter for more details).
- Proper provisioning of the VC capacities requires detailed knowledge of site-to-site traffic profiles, which are usually not readily available.


Last but not least, the overlay VPN model, when implemented with Layer 2 technologies, introduces another unnecessary layer of complexity into the New World Service Provider networks that are mostly IP-based, thus increasing the acquisition and operational costs of such a network.

Peer-to-peer VPN Model

The peer-to-peer VPN model was introduced a few years ago to alleviate the drawbacks of the overlay VPN model. In the peer-to-peer model, the Provider Edge (PE) device is a router (PE-router) that directly exchanges routing information with the CPE router. Figure 7-7 shows a sample peer-to-peer VPN, which is equivalent to the VPN in Figure 7-5.

NOTE

The *Managed Network* service offered by many service providers, where the service provider also manages the CPE devices, is not relevant to this discussion because it's only a repackaging of another service. The Managed Network provider concurrently assumes the role of the VPN service provider (providing the VPN infrastructure) and part of the VPN customer role (managing the CPE device).

Figure 7-7 *Sample Peer-to-peer VPN***NOTE**

Please note that this section describes the non-MPLS approach to peer-to-peer VPN as currently deployed by several large service providers and the complexities associated with it. The MPLS-based peer-to-peer VPN approach is described in the next chapter.

The peer-to-peer model provides a number of advantages over the traditional overlay model:

- Routing (from the customer's perspective) becomes exceedingly simple, as the customer router exchanges routing information with only one (or a few) PE-router, whereas in the overlay VPN network, the number of neighbor routers can grow to a large number.

- Routing between the customer sites is always optimal, as the provider routers know the customer's network topology and can thus establish optimum inter-site routing.
- Bandwidth provisioning is simpler because the customer has to specify only the inbound and outbound bandwidths for each site (Committed Access Rate [CAR] and Committed Delivery Rate [CDR]) and not the exact site-to-site traffic profile.
- The addition of a new site is simpler because the service provider provisions only an additional site and changes the configuration on the attached PE-router. Under the overlay VPN model, the service provider must provision a whole set of VCs leading from that site to other sites of the customer VPN.


Prior to an MPLS-based VPN implementation, two implementation options existed for the peer-to-peer VPN model:

- The shared-router approach, where several VPN customers share the same PE-router.
- The dedicated-router approach, where each VPN customer has dedicated PE-routers.

Shared-router Approach to Peer-to-peer VPN Model

In the shared-router approach, several customers can be connected to the same PE-router. Access lists have to be configured on every PE-CE interface on the PE-router to ensure isolation between VPN customers, to prevent a VPN customer from breaking into another VPN network, or to prevent a VPN customer from performing a denial-of-service attack on another VPN customer. Figure 7-8 illustrates a sample shared-router configuration.

Figure 7-8 Peer-to-peer VPN Model: Shared Router Configuration



Let's assume that the customers shown in Figure 7-8 use the address space and routing protocols from Table 7-1.

Table 7-1 *Peer-to-peer Shared-router Example—Address Space*

Customer Name	Address Space	Routing Protocol
FriedFoods (Customer #75)	155.13.0.0/16	RIP
GeneralMining (Customer #98)	195.166.16.0/20	OSPF (area 3)

To ensure the isolation between the customers, the configuration from Example 7-1 would have to be entered in the POP-router in Figure 7-8.

Example 7-1 *POP-router Configuration*


```

interface serial 0/0/1
  description FriedFoods - San Jose Site
  ip address 153.13.254.1 255.255.255.252
  ! The IP address on WAN link is an address from Customer's address space
  ip access-group FriedFoods in
  ip access-group FriedFoods out
  !
interface serial 0/0/2
  description FriedFoods - Santa Clara Site
  ip address 153.13.254.5 255.255.255.252
  ip access-group FriedFoods in
  ip access-group FriedFoods out
  !
interface serial 0/1/3
  description GeneralMining - Mountain View Site
  ip address 195.166.31.17 255.255.255.252
  ip access-group GeneralMining in
  ip access-group GeneralMining out
  !
router rip
  network 153.13.0.0
  !
router ospf 1
  network 195.166.31.17 0.0.0.0 area 3
  !
ip access-list FriedFoods
  permit ip 155.13.0.0 255.255.0.0 153.13.0.0 255.255.0.0
  !
ip access-list GeneralMining
  permit ip 195.166.16.0 255.255.240.0 195.166.16.0 255.255.240.0

```

Dedicated-router Approach to Peer-to-peer Model

In the dedicated-router approach to the peer-to-peer model, every VPN customer has their own dedicated PE routers (as detailed in Figure 7-9) and, thus, has access only to the routes contained within the routing table of that PE router.

Figure 7-9 Peer-to-peer VPN Model: Dedicated Router Configuration

The dedicated-router model uses routing protocols to create per-VPN routing tables on PE routers. The routing tables on PE-routers contain only the routes advertised by the VPN customer connected to them, resulting in almost perfect isolation between the VPN customers (assuming that the IP source routing is disabled). The routing in the dedicated-router model can be implemented as follows:

- Any routing protocol is run between the PE-router and the CE-router.
- BGP is run between the PE-router and the P-router.
- The PE-router redistributes routes received from the CE-router into BGP, marked with the customer ID (BGP community), and propagates the routes to the P-routers. P-routers thus contain all the routes from all VPN customers.
- P-routers propagate only routes with the proper BGP community to the PE-routers. The PE-routers thus receive only the routes that originated from the CE-routers in their VPN.

Relevant parts of PE-router and P-router configuration for the Service Provider Point-of-Presence (POP) shown in Figure 7-9 (assuming the address space and the routing protocols from Table 7-1) can be found in Example 7-2 and Example 7-3.

Example 7-2 PE-router Configuration

```
hostname PE-router-FriedFoods
!
interface serial 0/0/1
description FriedFoods - San Jose Site
ip address 153.13.254.1 255.255.255.252
```

continues

Example 7-2 *PE-router Configuration (Continued)*

```

!
interface serial 0/0/2
  description FriedFoods - Santa Clara Site
  ip address 153.13.254.5 255.255.255.252
!
interface FastEthernet 2/0/0
  description Intra-POP LAN
  ip address 10.13.1.2 255.255.255.0
!
router rip
  network 153.13.254.1
  version 2
  redistribute bgp 111 subnets
!
router bgp 111
  no auto-summary
  redistribute rip route-map ToBGP-FriedFoods
  neighbor 10.13.1.1 remote-as 111
!
route-map ToBGP-FriedFoods permit 10
  set community 111:75

```

Example 7-3 *P-router Configuration*

```

hostname P-Router-Silicon-Valley-POP
!
interface FastEthernet 0/1/0
  description Intra-POP LAN
  ip address 10.13.1.1 255.255.255.0
!
router bgp 111
  neighbor 10.13.1.2 remote-as 111
  neighbor 10.13.1.2 route-reflector-client
  neighbor 10.13.1.2 route-map VPN-FriedFoods out
!
route-map VPN-FriedFoods permit 10
  match community-list 75
!
ip community-list 75 permit 111:75

```

Comparison of Peer-to-peer Models

As you easily can deduce from Example 7-1, the shared-router peer-to-peer model is very hard to maintain because it requires the deployment of potentially long and complex access lists on almost every router interface. The dedicated-router approach, although simpler to configure and maintain, becomes very expensive for the service provider when it tries to serve a large number of customers with geographically dispersed sites.

Both peer-to-peer models also share several common drawbacks that prevent their widespread usage:

- All the customers share the same IP address space, preventing the customers from deploying private IP addresses according to RFC 1918. The customers must use either public IP addresses or private IP addresses allocated to them by the service provider.
- The customers cannot insert the default route into their VPN. This limitation prevents certain routing optimizations and prevents the customers from getting Internet access from another service provider.

In addition to these two drawbacks, the shared-router model suffers from additional complexity when several customers use the routing protocols (RIP, RIPv2, BGP, and IS-IS) where multiple instances are not supported in Cisco IOS.

Typical VPN Network Topologies

The VPN topology required by an organization should be dictated by the business problems the organization is trying to solve. However, several well-known topologies appear so often that they deserve to be discussed here. As you can see, the same topologies solve a variety of different business issues in different vertical markets or industries.

The VPN topologies discussed here can be split into three major categories:

- Topologies influenced by the overlay VPN model, which include hub-and-spoke topology, partial or full-mesh topology, and hybrid topology.
- Extranet topologies, which include any-to-any Extranet and Central Services Extranet.
- Special-purpose topologies, such as VPDN backbone and Managed Network topology.

Hub-and-spoke Topology

The most commonly encountered topology is a hub-and-spoke topology, where a number of remote offices (spokes) are connected to a central site (hub), similar to the setup in Figure 7-10. The remote offices usually can exchange data (there are no explicit security restrictions on inter-office traffic), but the amount of data exchanged between them is negligible. The hub-and-spoke topology is used typically in organizations with strict hierarchical structures, for example, banks, governments, retail stores, international organizations with small in-country offices, and so on.

NOTE

When deploying VPNs based on Layer 2 technologies, such as Frame Relay or ATM, the hub-and-spoke VPN topology is more common than you might expect. This is based purely on business needs due to higher costs or increased routing complexity associated with other topologies that use these types of technologies. In other words, there are many examples where the customer could benefit from a different topology but has nonetheless chosen the hub-and-spoke topology for cost or complexity reasons.

With increased redundancy requirements, the simple hub-and-spoke topology from Figure 7-10 often is enhanced with an additional router at the central site (shown in Figure 7-11) or with a backup central site, which is then linked with the primary central site through a higher-speed connection (shown in Figure 7-12).

Figure 7-10 *Hub-and-spoke Topology*






Figure 7-11 *Hub-and-spoke Topology with Two Central Routers*



Implementing redundant hub-and-spoke topology with an overlay VC-based VPN model always poses a number of challenges. Each hub site requires a VC to at least two central routers. These VCs could be provisioned in primary-backup configuration or in load-sharing configuration with a number of drawbacks of one or the other solution:

- In primary-backup configuration, the backup VC is unused while the primary VC is active, resulting in unnecessary expenses incurred by the customer.
- In load-sharing configuration, the spoke site encounters reduced throughput if one of the VCs (or one of the central routers) fails. The load-sharing configuration is also not appropriate for the topologies with a backup central site similar to the one in Figure 7-12.



Figure 7-12 Hub-and-spoke Topology with Two Central Sites



The higher-quality service providers try to meet the redundancy requirements of their customers with an enhanced service offering called *shadow PVC*. With a shadow PVC, the customer gets two virtual circuits for the price of one on the condition that they can use only one VC for data traffic at a time (a small amount of traffic is allowed on the second PVC to enable routing protocol exchanges over the second PVC).

Redundancy requirements can further complicate hub-and-spoke topology with the introduction of dial-backup features. The dial backup solution implemented within the service provider network (for example, an ISDN connection backing up a Frame-Relay leased line, as shown in Figure 7-13) is transparent to the customer, but it does not offer true end-to-end redundancy because it cannot detect all potential failures (for example, CPE or routing protocol failures). The true end-to-end redundancy in an overlay VPN model can be achieved only by CPE devices establishing a dial-up connection outside the VPN space.

Usually, simple hub-and-spoke topology transforms into multilevel topology as the network grows. The multilevel topology can be a recursive hub-and-spoke topology, similar to the one shown in Figure 7-14, or a hybrid topology, which is discussed later in this section. The network restructuring can be triggered by scalability restrictions of IP routing protocols or by application-level scalability issues (for example, the introduction of a three-tier client-server approach).

Figure 7-13 *Dial Backup Solution Within Service Provider Network***Figure 7-14** *Multilevel Hub-and-spoke Topology*

The hub-and-spoke topology implemented with an overlay VPN model is well suited to environments where the remote offices mostly exchange data with the central sites and not with each other, as the data exchanged between the remote offices always gets transported via the central site. If the amount of data exchanged between the remote offices represents a significant proportion of the overall network traffic, partial-mesh or full-mesh topology might be more appropriate.

Partial- or Full-mesh Topology

Not all customers can implement their networks with the hub-and-spoke topology discussed in the previous section for a variety of reasons, for example:

- The organization might be less hierarchical in structure, requiring data exchange between various points in the organization.
- The applications used in the organization need peer-to-peer communication (for example, messaging or collaboration systems).
- For some multinational corporations, the cost of hub-and-spoke topology might be excessive due to the high cost of international links.

In these cases, the overlay VPN model best suited to the organization's needs would be a partial-mesh model, where the sites in the VPN are connected by VCs dictated by traffic requirements (which eventually are dictated by business needs). If not all sites have direct connectivity to all other sites (like the example in Figure 7-15), the topology is called a *partial mesh*; if every site has a direct connection to every other site, the topology is called a *full mesh*.

NOTE

Not many full-mesh networks are implemented due to the very high cost of this approach and the complexity introduced by the high number of VCs. With this type of topology, the number of VCs = $[(n-1) \times n] \div 2$ where n is equal to the number of attached devices.

Most of the customers have to settle for a partial mesh topology, which usually is affected by compromises and external parameters, such as link availability and the cost of VCs.

W

WANS

ATM

- Cell-mode MPLS operation, 69
- heterogeneous MPLS-mode operation, 69
- MPLS operation, 67, 69

core redundancy, 7

Frame Relay, MPLS operation, 65–67

Layer 2 paths, selecting, 6–7

virtual circuits, provisioning, 8

Web sites, IETF, 11, 177

MPLS AND VPN ARCHITECTURES

LIBRARY OF CONGRESS



- In-depth analysis of the Multiprotocol Label Switching (MPLS) architecture
- Detailed discussion of the mechanisms and features that constitute the architecture
- Learn how MPLS scales to support tens of thousands of Virtual Private Networks (VPNs)
- Extensive case studies guide you through the design and deployment of real-world MPLS/VPN networks
- Configuration examples and guidelines assist in configuring MPLS on Cisco devices
- Design and implementation options help you build various VPN topologies

This book is part of the Cisco Press Core Series, which offers networking professionals valuable information for constructing efficient networks, understanding new technologies, and building successful careers.

Category: Networking—Design
Covers: MPLS and VPNs

CISCO SYSTEMS

CISCO PRESS
ciscopress.com

Multiprotocol Label Switching (MPLS) is a high-performance packet forwarding technology, both within a service provider's enterprise network, and the most widely deployed usage today is the enabling of Virtual Private Networks (VPNs). With the introduction of MPLS-enabled VPNs, network designers are able to better scale their networks than with the methods available in the past.

Network engineers and administrators need quick, effective education on this technology to efficiently deploy MPLS-enabled VPNs within their networks. With that goal in mind, *MPLS and VPN Architectures* provides the first in-depth discussion particular to Cisco's MPLS architecture. This book covers MPLS theory and configuration, network design issues, and case studies as well as one major MPLS application: MPLS-based VPNs. The MPLS/VPN architecture and all of its mechanisms are explained with configuration examples, suggested design and deployment guidelines, and extensive case studies.

MPLS and VPN Architectures is your practical guide to understanding, designing, and deploying MPLS and MPLS-based VPNs.



Ivan Pepelnjak (CCIE™ #1354) has more than 10 years experience in designing, installing, troubleshooting, and operating large corporate and service provider WANs and LANs. As technical director of NIL Communications, his expertise has benefited such service providers as British Telecom, Deutsche Telekom, France Telecom, BelgaCom, Unisource, Cegetel, and Siris. Ivan is one of the top Cisco® routing authorities in Europe and is the author of *EIGRP Network Design Solutions*.



Jim Guichard (CCIE #2069) is a Network Design Consultant at Cisco Systems® who has extensive experience with designing, migrating, and deploying large-scale IP networks. Jim has used his practical knowledge to assist many of Cisco's largest customers with their network rollouts. He has worked closely with Cisco development and has spent much time researching and testing MPLS/VPN functionality, both within the lab environment and within multiple customer deployments throughout Europe.

\$55.00 USA
\$81.95 CAN




6 19472 05002 6

ISBN 1-58705-002-1



9 781587 050022



5 5500

MPIS AND VPN ARCHITECTURE

TK 5105

.573

.G85

2001

V. 1

Copy 1